

DOCKET NO. RSW920010072US1

#14
@B
4/2/04

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant: M. Trevathan

Examiner: Choi, Woo H.

Serial No.: 09/826,085

Art Unit: 2186

Filed: 4/4/2001

For: **IMPROVED METHOD FOR MANAGING A CACHE MEMORY USING A
PREDICTIVE MODELING ENGINE TO SELECT A CACHING ALGORITHM**

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

RECEIVED

MAR 31 2004

Technology Center 2100

REPLY BRIEF OF APPELLANT

This Reply Brief addresses issues raised in the Examiner's Answer mailed February 24, 2004.

Issue 1

CLAIMS 1, 3-7, AND 9 ARE NOT ANTICIPATED BY GAITHER (US PATENT NO. 6,223,256).

The Examiner rejected claims 1, 3-7 and 9 under 35 U.S.C. §102(e) as allegedly being anticipated by Gaither (US Patent No. 6,223,256).

Claim 1

Appellant respectfully contends that the Examiner's Answer does not adequately support the Examiner's contention that Gaither discloses selecting a preferred caching algorithm **in response to** the act of analyzing information stored in a caching profile, as required by claim 1.

In particular, the Examiner's Answer cites Gaither's abstract as allegedly disclosing that the

09/559,710

Admitted
4/30/04
wlc

“cache class attribute is used to select a replacement algorithm” (see Examiner’s Answer, page 4, lines 2-3). Appellant contends, however, that the act of analyzing information stored in a caching profile does not trigger selecting a preferred caching algorithm. Indeed, the act of analyzing information stored in a caching profile is performed only **periodically**. See Gaither, col.8, lines 42-45 (“the cache may be **periodically** observed to detect entries that are frequent or infrequent, entry residency time, and replacement history information” (emphasis added)). In contrast, the selection of a preferred caching algorithm occurs **every time** real-time data is ready to be placed in the cache. **Gaither does not disclose any correlation between when the caching profile is analyzed and when the preferred caching algorithm is selected.** Therefore, the act of analyzing information stored in a caching profile does not trigger selecting a preferred caching algorithm, which means that the selection of a preferred caching algorithm is not in response to the act of analyzing information stored in a caching profile.

Gaither teaches that every time real-time data is to be stored in the cache, the cache replacement algorithm is selected based on the class attribute of both the real-time data to next be stored in the cache and of the data currently stored in the cache, as illustrated by Gaither in numerous examples such as the examples in col. 9, line 36 - col. 10, line 21. In fact, because the cache replacement algorithm in Gaither is a function of the class attribute of the real-time data to be next stored in the cache, a cache replacement algorithm must be selected every time real-time data is to be stored in the cache. Therefore, the selection of the cache replacement algorithm is triggered by the need to store the real-time data in the cache and not by the analysis of the information stored in a caching profile which occurs only periodically. Appellant reiterates that such analysis is not disclosed by Gaither as having any correlation with when real-time data is to

be stored in the cache. Even though such analysis may result in updating class attributes of data that relates to the information stored in the caching profile, and even though the selection of the cache replacement algorithm is selected based on the class attributes, the actual act of selecting the cache replacement algorithm is not triggered by the updating of the class attributes but is instead triggered by the need to store real-time data in the cache. Therefore, the selection of a preferred caching algorithm is not in response to the act of analyzing information stored in a caching profile as required by claim 1.

Based on the preceding arguments, Appellant respectfully maintains that Gaither does not anticipate claim 1, and that claim 1 is in condition for allowance. Accordingly, Appellant contends that the rejection of claim 1 under 35 U.S.C. §102(e) is improper and should be reversed.

Claim 3

Appellant respectfully contends that Gaither does not anticipate claim 3, because Gaither does not teach each and every feature of claim 3. For example, Gaither does not teach the feature: “**responsive to the act of updating [the caching profile], analyzing information stored in the caching profile**” (emphasis added). The Examiner’s Answer argues: “responsive to the acts of updating, analyzing information stored in the caching profile (as noted above, information collection and analyses are done dynamically. Analyzing information is necessarily responsive to the act of updating, i.e. information gathering, since the analysis cannot be done without the information to be analyzed. Dynamic analysis implies that the analysis is done on the fly’ as the information is updated)”. Examiner’s Answer, page 5, lines 3-7.

Appellant contends the preceding argument by the Examiner , namely that “the analysis cannot be done without the information to be analyzed”, skirts the issue of causality which is inherent in the preceding feature of claim 3. The issue is one of cause and effect; i.e. whether updating the caching profile **triggers** the analysis of information stored in the caching profile. Appellant contends that Gaither answers this question as follows: “the cache may be **periodically** observed to detect entries that are frequent or infrequent, entry residency time, and replacement history information (emphasis added)”. Gaither, col.8, lines 42-45. In other words, the key to this issue is the word “periodically” and Gaither discloses that the trigger that drives the analysis of information stored in the caching profile is the periodic passage of time which is independent of updating the caching profile. For example, let’s assume that the periodic observing and detection occurs at 10-second time intervals. In any given 10-second time interval, the caching profile may have been updated many times or may not have been updated at all. The point is that the analysis of information stored in the caching profile is being triggered by the periodic passage of time on the clock and not by the updating of the caching profile. Appellant contends that the only trigger disclosed by Gaither for analyzing information stored in the caching profile is the periodic passage of time. Gaither does not disclose any other trigger for analyzing information stored in the caching profile. Gaither most certainly does not disclose anywhere that the act of updating the caching profile triggers the analysis of information stored in the caching profile.

To emphasize the point, even if Gaither discloses both updating the caching profile and analyzing information stored in the caching profile, the issue is one of cause in effect. Does the alleged updating the caching profile **trigger** the alleged analysis of information stored in the caching profile in Gaither? As explained *supra*, the answer is most assuredly NO, since only the

periodic passage of time is disclosed by Gaither as triggering the analysis of information stored in the caching profile.

Based on the preceding arguments, Appellant respectfully maintains that Gaither does not anticipate claim 3, and that claim 3 is in condition for allowance. Accordingly, Appellant contends that the rejection of claim 3 under 35 U.S.C. §102(e) is improper and should be reversed.

Claim 5

Appellant respectfully contends that Gaither does not anticipate claim 3, because Gaither does not teach each and every feature of claim 3. For example, Gaither not teach the feature: “**responsive to** arrival of a file at a cache ..., analyzing information stored in a caching profile ...” (emphasis added). The Examiner’s Answer argues: “responsive to arrival of a file at a cache (see the discussion of arrival of a compiled program above), analyzing information stored in a caching profile (i.e. dynamic analysis of gathered information) ...”. Examiner’s Answer, page 5, lines 14-16. As to the discussion of arrival of a compiled program, the Examiner argues: “A program that is compiled is stored in one or more files. Run-time caching behavior information of a program cannot be collected until the program, i.e. file, arrives at a computer memory and more specifically at a cache. The collection of information gathered constitutes a ‘caching profile’ since it contains caching behavior information. Gathering of information requires constant updating of this ‘caching profile’ as new information is generated and gathered while the program runs”. Examiner’s Answer, page 4, line 17 - page 5, line 2.

Appellant notes that the “responsive to” issue in claim 5 is analogous to the “responsive

to” issue in claim 3. As with claim 3, the preceding argument by the Examiner with respect to claim 5 again skirts the issue of causality which is inherent in the preceding feature of claim 5. The issue is one of cause and effect; i.e. whether arrival of a file at a cache **triggers** the analysis of information stored in the caching profile. Appellant contends that Gaither answers this question as follows: “the cache may be **periodically** observed to detect entries that are frequent or infrequent, entry residency time, and replacement history information (emphasis added)”. Gaither, col.8, lines 42-45. In other words, the key to this issue is the word “periodically” and Gaither discloses that the trigger that drives the analysis of information stored in the caching profile is passage of time which is independent of arrival of a file at a cache. For example, let’s assume that the periodic observing and detection occurs at 10-second time intervals. In any given 10-second time interval, the caching profile may have been updated many times or may not have been updated at all. The point is that the analysis of information stored in the caching profile is being triggered by passage of time on the clock and not by the arrival of a file at a cache. Appellant contends that the only trigger disclosed by Gaither for analyzing information stored in the caching profile is the periodic passage of time. Gaither does not disclose any other trigger for analyzing information stored in the caching profile. Gaither most certainly does not disclose anywhere that the arrival of a file at a cache triggers the analysis of information stored in the caching profile.

To emphasize the point, even if Gaither discloses both arrival of a file at a cache and analyzing information stored in the caching profile, the issue is one of cause in effect. Does the alleged arrival of a file at a cache trigger the alleged analysis of information stored in the caching profile in Gaither? As explained *supra*, the answer is most assuredly: NO, since only the

periodic passage of time is disclosed by Gaither as triggering the analysis of information stored in the caching profile.

In addition, Appellant respectfully contends that Gaither does not anticipate claim 5, because Gaither does not teach “responsive to a comparison of the metrics [of the plurality of metrics] one with another, selecting a preferred caching algorithm from a plurality of caching algorithms”.

The Examiner’s Answer argues that “responsive to a comparison of the metrics one with another, selecting preferred caching algorithm from a plurality of caching algorithms” Examiner’s Answer, page 6, lines 3-4. The Examiner considers metrics to include “frequency and infrequency of entries, residency time, and cache history.” Examiner’s Answer, page 5, lines 19 - page 6, line 1.

Appellant contends that the Examiners argument is not persuasive because in Gaither the preferred caching algorithm is not based on comparing frequency and infrequency of entries, residency time, and cache history. Indeed, Gaither determines a preferred caching algorithm based on utilizing the known class attribute of dynamic information ready to be cached and the known class attribute(s) of the information currently stored in the cache as discussed extensively in Gaither, including in col. 9, line 36 - col. 10, line 21. The preferred caching algorithm is determined on the basis of an algorithm that utilizes aforementioned class attributes. Although the metrics of frequency and infrequency of entries, residency time, and cache history may be used to update class attributes, it is not possible in Gaither to determine the preferred caching algorithm from knowledge of the metrics of frequency and infrequency of entries, residency time,

and cache history as the Examiner's Answer alleges. Accordingly, Appellant contends that the Examiner has failed to establish a *prima facie* case of obviousness in relation to claim 5.

Based on the preceding arguments, Appellant respectfully maintains that Gaither does not anticipate claim 5, and that claim 5 is in condition for allowance. Accordingly, Appellant contends that the rejection of claim 5 under 35 U.S.C. §102(e) is improper and should be reversed.

Claim 6

Since claim 6 depends from claim 5, which Appellant has argued *supra* to be patentable and in condition for allowance under 35 U.S.C. §102(e), Appellant maintains that claim 6 is likewise patentable and in condition for allowance.

In addition with respect to claim 6, Appellant respectfully contends that Gaither does not teach "wherein the plurality of metrics includes clustering metrics". The Examiner's Answer alleges "[Gaither teaches] ... the plurality of metrics includes clustering metrics ... (col. 13, line 4, non-uniform distribution implies clustering)." Examiner's Answer, page 6, lines 6-7.

In response, Appellant contends that the Examiner's argument does not distinguish between cause and effect as required by claim 6, as will be next explained. Appellant notes that col. 13, line 4 of Gaither is referring to "hierarchical sections with non-uniform distributions", which Gaither explains as follows in col. 12, lines 51-64 in relation to FIG. 4:

"In FIG. 4, the horizontal axis represents locations in a cache numbered from 0-N. The

vertical axis represents the probability that a replacement algorithm will place an item of a particular class in a particular location. For items having Class A attributes, the probability of being placed in the locations near location 0 is higher than the probability of being placed in the locations near location N. For items having Class B attributes, the probability of being placed in the locations near location N is higher than the probability of being placed in the locations near location 0. Therefore, each class has access to the entire cache, but on a statistical basis, victimization of one class by the other class is reduced because the spatial distributions are not uniform.”

From the preceding explanation by Gaither, it is clear that the “non-uniform distribution” in Gaither pertains to a situation in which the probability that an item will be placed in a location I in the cache ($I=0, 1, 2, \dots, N$) is a function of I rather than being independent of I . Therefore the metric associated with this “non-uniform distribution” is not a metric that claim 6 reads on, because this “non-uniform distribution” results from a selection of a caching algorithm, as taught by Gaither on col. 12, lines 46-50 (“with multiple replacement algorithms, different classes can be spatially distributed differently within one or more sections, changing the probability of victimization without distinct section boundaries”). However, claim 6 requires (though its dependence on claim 5) that the selection of a caching algorithm results from a comparison of metrics with one another, which is logically impossible if the “non-uniform distribution” is a consequence of a choice of caching algorithm. In other words, claim 6 requires that a comparison of metrics is a cause of the selected caching algorithm, whereas the non-uniform distribution in Gaither (which the Examiner alleges to be a clustering metric of claim 6) is not a

cause of the selected caching algorithm, but is rather a consequence of the selected caching algorithm. In summary, Appellant contends that the Examiner's argument in relation to claim 6 is not persuasive, because the Examiner's failure to consider the cause and effect relationships discussed *supra* has resulted in the Examiner incorrectly concluding that Gaither's non-uniform distribution is a metric recited in claim 6.

The Examiner's Answer has responded nonpersuasively to the preceding argument by Appellant that also appeared in Appellant's appeal brief. See Examiner's Answer, page 15, line 22 - page 16, line 16. Appellant notes that the preceding response in the Examiner's Answer has not invalidated Appellant's argument that it is logically impossible for the non-uniform distribution in Gaither to be both a cause and a consequence of the caching algorithm. In fact, the response in the Examiner's Answer has not even attempted to demonstrate a flaw in Appellant's argument. The preceding logical inconsistency uncovered by Appellant results from the Examiner incorrectly concluding that Gaither's non-uniform distribution is a clustering metric that determines the preferred caching algorithm.

Furthermore, the Examiner's argument with respect to how Appellant's specification describes a "metric" is irrelevant to the cause and effect issue and has no bearing on Appellant's assertion that the non-uniform distribution in Gaither cannot logically be both a cause and a consequence of the caching algorithm. In effect, the Examiner's argument that Gaither's non-uniform distribution implies that the plurality of metrics includes clustering metrics cannot be correct, since it is logically inconsistent.

Based on the preceding arguments, Appellant respectfully maintains that Gaither does not anticipate claim 6, and that claim 6 is in condition for allowance. Accordingly,

Appellant contends that the rejection of claim 6 under 35 U.S.C. §102(e) is improper and should be reversed.

Claim 7

Since claim 7 depends from claim 5, which Appellant has argued *supra* to be patentable and in condition for allowance under 35 U.S.C. §102(e), Appellant maintains that claim 6 is likewise patentable and in condition for allowance.

In addition with respect to claim 7, Appellant respectfully contends that Gaither does not teach “wherein the plurality of metrics includes scattering metrics”. The Examiner’s Answer alleges: “[Gaither teaches] ... the plurality of metrics includes scattering metrics (col. 13, line 3, in a uniform distribution things are scattered evenly).”

In response, Appellant contends that the examiner’s argument does not distinguish between cause and effect as required by claim 7, as will be next explained. Appellant notes that col. 13, line 3 of Gaither is referring to “hierarchical sections with uniform distributions”, which are the opposite of “non-uniform distributions”. Gaither explains non-uniform distributions as follows in col. 12, lines 51-64 in relation to FIG. 4:

“In FIG. 4, the horizontal axis represents locations in a cache numbered from 0-N. The vertical axis represents the probability that a replacement algorithm will place an item of a particular class in a particular location. For items having Class A attributes, the probability of being placed in the locations near location 0 is higher than the probability of being placed in the locations near location N. For items having Class B attributes, the probability of being placed in the locations near location N is higher than the probability

of being placed in the locations near location 0. Therefore, each class has access to the entire cache, but on a statistical basis, victimization of one class by the other class is reduced because the spatial distributions are not uniform.”

From the preceding explanation by Gaither, it is clear that the “uniform distribution, which is the opposite of the “non-uniform distribution”, pertains to a situation in which the probability that an item will be placed in a location I in the cache ($I=0, 1, 2, \dots, N$) is independent of I . Therefore the metric associated with this “uniform distribution” is not a metric that claim 7 reads on, because this “uniform distribution” results from a selection of a caching algorithm, as taught by Gaither on col. 12, lines 46-50 (“with multiple replacement algorithms, different classes can be spatially distributed differently within one or more sections, changing the probability of victimization without distinct section boundaries”). However, claim 7 requires (though its dependence on claim 5) that the selection of a caching algorithm results from a comparison of metrics with one another, which is logically impossible if the “uniform distribution” is a consequence of a choice of caching algorithm. In other words, claim 7 requires that a comparison of metrics is a cause of the selected caching algorithm, whereas the uniform distribution in Gaither (which the Examiner alleges to be a scattering metric of claim 7) is not a cause of the selected caching algorithm, but is rather a consequence of the selected caching algorithm. In summary, Appellant contends that the Examiner’s argument in relation to claim 7 is not persuasive, because the Examiner’s failure to consider the cause and effect relationships discussed *supra* has resulted in the Examiner incorrectly concluding that Gaither’s uniform distribution is a metric of claim 7.

The Examiner’s Answer has responded nonpersuasively to the preceding argument by

Appellant that also appeared in Appellant's appeal brief. See Examiner's Answer, page 15, line 22 - page 16, line 16. Appellant notes that the response in the Examiner's Answer has not invalidated Appellants argument that it is logically impossible for the uniform distribution in Gaither to be both a cause and a consequence of the caching algorithm. In fact, the response in the Examiner's Answer has not even attempted to demonstrate a flaw in Appellant's argument. The preceding logical inconsistency uncovered by Appellant results from the Examiner incorrectly concluding that Gaither's uniform distribution is a scattering metric that determines the preferred caching algorithm.

Furthermore, the Examiner's argument with respect to how Appellant's specification describes a "metric" is irrelevant to the cause and effect issue and has no bearing on Appellant's assertion that the non-uniform distribution in Gaither cannot logically be both a cause and a consequence of the caching algorithm. In effect, the Examiner's argument that Gaither's uniform distribution implies that the plurality of metrics includes scattering metrics cannot be correct, since it is logically inconsistent.

Based on the preceding arguments, Appellant respectfully maintains that Gaither does not anticipate claim 7, and that claim 7 is in condition for allowance. Accordingly, Appellant contends that the rejection of claim 7 under 35 U.S.C. §102(e) is improper and should be reversed.

Issue 2

CLAIM 8 IS NOT UNPATENTABLE UNDER 35 U.S.C. §103(a) OVER ARLITT ET AL. (US PATENT NO. 6,272,598) IN VIEW OF ADMITTED PRIOR ART.

The Examiner rejected claim 8 under 35 U.S.C. §103(a) as allegedly being unpatentable over Arlitt et al. (US Patent No. 6,272,598) in view of admitted prior art.

As a first example illustrating that Arlitt in view of the admitted prior art does not teach or suggest features of claim 8, Appellant contends that Arlitt in view of the admitted prior art does not teach or suggest the feature: “responsive to arrival of a file at a cache, analyzing information stored in a caching profile by computing a plurality of metrics”.

The Examiner’s Answer argues that “Arlitt discloses ... responsive to arrival of a file at a cache, analyzing information stored in a caching profile by computing a plurality of metrics (col. 5, lines 48- 50, **arrival** of a file is inherent since the subject matter of Arlitt's invention is a Web cache which caches Web page documents, see col. 1, lines 15 - 29. Hit rate and byte hit rate computations are necessarily responsive to arrival of a file. A hit rate computation can only begin in response to arrival of a file, since the reference to the file prior to its arrival is classified a ‘miss’, see col. 1, lines 53 - 54. Determination of hit rate and byte hit rate, i.e. ‘cache profile’, requires computations and storage.)”. Examiner’s Answer, page 7, lines 7-14.

The Examiner’s Answer clarifies its argument on page 18, lines 12-22 as follows: “
“arrival of a file at a cache” is taught since this is inherent in Arlitt's invention. The subject matter of Arlitt's invention is Web caching (see title). Arlitt's Web cache caches Web documents (see col. 1, lines 15 - 29) that are composed of various files. Each of the files is referred to as an “object” in Arlitt's disclosure (col. 1, lines 23 -24). The Web cache caches these objects. Web page files must arrive at and reside in the cache in order for Arlitt's invention to be functional. Secondly, “responsive

to arrival .. ." is taught as well. In the art of caching, a cache hit occurs when a requested object is present in the cache. A cache miss is declared when the requested object is not in the cache (see col. 1, lines 53 - 54). Hit rates cited above cannot be determined in the absence of an object or file in the cache. Determination of hit rates can only be conducted in response to arrival of a file, since a reference to the file prior to its arrival is classified as a "miss"."

In response, Appellant submits that the Examiner's Answer has misinterpreted Arlitt and has incorrectly described the circumstances under which Arlitt teaches the computation of a "hit rate" and a "miss rate". Arlitt describes this process in col. 1, line 48 - col. 2, line 9 with respect to FIG. 1 in Arlitt, and also in col. 4, lines 15-27 of Arlitt. The cited description by Arlitt in col. 1, line 48 - col. 2, line 9 is as follows:

"As can be seen from FIG. 1, when a client device or user terminal 12 generates a request for a particular object (e.g., the object 10 stored in the remote server 18), the cache of the proxy server 16 in the local server 14 is searched to determine whether the object 10 is stored at the proxy server 16. If the object is not found in the cache of the proxy server 16, a "cache miss" results and the local server 14 directs the request to the remote server 18 via the Internet 20. As can be seen from FIG. 1, the remote server 18 stores the requested object 10. Once the remote server 18 receives the request, it directs a response with the requested object 10 to the client device 12 via the local server 14. During this process, the requested object 10 is also cached in the proxy server 16 of the local server 14. This eliminates the need for the local server 14 to send another request to the remote server 18 for the same object 10 at a later time when either the same client device 12 or a different client device (not shown) requests for the same object 10. When the object 10 is again

requested, the proxy server 16 is accessed and a "cache hit" results. In this case, the cached object 10 is quickly forwarded to the client device directly from the proxy server 16. This eliminates delays encountered in communicating between the proxy server 16 and the remote server 18. By storing copies of objects received from remote sites, the proxy server 16 reduces the number of requests that are directed to the remote server 18, as well as the traffic on the Internet 20 as a result of transmitting the responses in the form of a number of packets that must be reassembled at the client device 12."

Appellant submits that the preceding description by Arlitt makes it clear that the "hit rates" and "miss rates" are capable of being updated after the cache of the local server is searched in response to a request for an object (e.g., web page) from a client device and well before the web page is ever transmitted from the remote server to the local server. All that the local server needs to know to update the "hit rate" and "miss rate" is whether the requested object is currently stored in the cache of the local server and arrival of a the web page at the cache of the local server is thus not needed to update the "hit rate" and "miss rate". In fact, when a "hit" occurs, the web page is not even transmitted from the remote server and therefore does not "arrive" at the cache of the local server.

In essence, updating the "hit rate" and "miss rate" in Arlitt is responsive to a message received at the local server from the client device. The message is received at the local server (in the form of packets) and may trigger analysis of information stored in the cache in order to determine whether a hit or miss has occurred. Although the message is received at local server, the message is not received at the cache of the local server as required by claim 8. Therefore,

Arlitt does not teach or suggest “responsive to arrival of a file at a cache, analyzing information stored in a caching profile by computing a plurality of metrics”. Accordingly, claim 8 is not unpatentable over Arlitt in view of the admitted prior art.

As a second example illustrating that Arlitt in view of the admitted prior art does not teach or suggest other features of claim 8, Appellant contends that Arlitt in view of the admitted prior art does not teach or suggest the feature: “responsive to a comparison of the metrics one with another, selecting a preferred caching algorithm from a plurality of caching algorithms”.

The Examiner’s Answer argues that “Arlitt discloses ... responsive to a comparison of the metrics (i.e., the hit rate and byte rate) one with another, selecting preferred caching algorithm from a plurality of caching algorithms (col. lines 35 - 47, Arlitt discloses that "Each replacement policy is optimized for one performance metric." Arlitt makes it clear that the use of a particular replacement policy, i.e. caching algorithm, is based on the performance metric. Algorithm selections based on the performance metrics require comparison of metrics).”

In response to the preceding argument by the Examiner, Appellant notes that Arlitt does not disclose in col. 5, lines 35-47 of Arlitt that said selection of a caching algorithm is “responsive to a comparison of the metrics one with another” as required by claim 8. The Examiner’s Answer acknowledges that Arlitt does not explicitly disclose a comparison of metrics but nonetheless argues: “The passage in Arlitt reference cited in the rejection, col. 5, lines 35 - 47, discloses that the cache manager uses multiple replacement strategies (caching algorithms). Arlitt specifically states that "Each **replacement policy** is optimized for one **performance metric**." Arlitt makes it clear that the use of a particular replacement policy (caching algorithm) is based on

the performance metric. How would a cache manager select a replacement strategy from a number of replacement strategies, each of which are keyed to one performance metric, without examining and comparing the associated metrics when they are the bases for the selection of the replacement strategy to be used? The metrics must be compared in order for Arlitt's invention to function as described.”. Examiner’s Answer, page 19, lines 9-17.

In response to the preceding argument in the Examiner’s Answer, Appellant contends that the preceding analysis by the Examiner is unsupported speculation. Appellant next offers a more accurate analysis of what Arlitt actually discloses. Arlitt discloses that the cache is subdivided into discrete storage areas and each such storage area stores one particular type of object such as image objects, video objects, large objects, extremely popular objects, etc. (see Arlitt, col. 3, lines 45-52, 62-67). Each such storage area has one replacement policy which the storage area receives from a cache manager (see Arlitt, col. 7, lines 36-40; col. 8, lines 55-59). Thus, when an object is to be stored in the cache, the object’s type is determined and the object is then sent to the storage area of the cache that has been designated for that type of the object. The object is then stored in that storage area in accordance with the replacement policy that has been assigned to that storage area.

In Appellant’s preceding analysis, the key unanswered question is this: how does the cache manager know what cache replacement policy to assign to each storage area? Unfortunately, Arlitt does not answer this key question. The Examiner’s Answer has speculated that this assignment must be made on the basis of comparison of metrics. However, Appellant can speculate that this assignment is based on the type of object that the storage area stores. In reality, the answer to this question not reasonably known by anyone - not by Examiner and not by

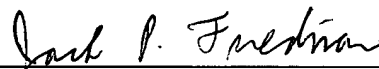
Appellant. Therefore, Appellant concludes that Arlitt does not disclose "responsive to a comparison of the metrics one with another, selecting a preferred caching algorithm from a plurality of caching algorithms" as required by claim 8. Appellant additionally concludes that Arlitt's failure to answer this key question means that the preceding feature of claim 8 is not enabled by Arlitt which invalidates the use of Arlitt as a reference.

Based on the preceding arguments, Appellant respectfully maintains claim 8 is not unpatentable over Arlitt in view of the admitted prior art, and that claim 8 is in condition for allowance.

SUMMARY

In summary, Appellant respectfully requests reversal of the July 17, 2003 office action rejection of claims 1 and 3-9.

Respectfully submitted,

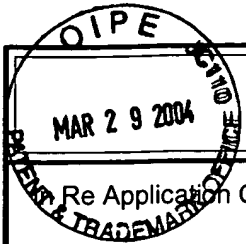


Jack P. Friedman
Attorney For Appellant
Registration No. 44,688

Dated: 03/25/2004

Schmeiser, Olsen & Watts
3 Lear Jet Lane - Suite 201
Latham, New York 12110
(518) 220-1850

AF/2702



TRANSMITTAL LETTER
(General - Patent Pending)

Docket No.
RSW920010072US1

Re Application Of: **M. Trevathan**

Serial No.
09/826,085

Filing Date
4/4/2001

Examiner
Choi, Woo H.

Group Art Unit
2186

Title: **IMPROVED METHOD FOR MANAGING A CACHE MEMORY USING A PREDICTIVE MODELING ENGINE TO SELECT A CACHING ALGORITHM**

TO THE COMMISSIONER FOR PATENTS:

RECEIVED

MAR 31 2004

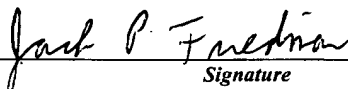
Technology Center 2100

Transmitted herewith is:

Reply Brief (19 pages)

in the above identified application.

- ☒ No additional fee is required.
- ☐ A check in the amount of _____ is attached.
- ☒ The Director is hereby authorized to charge and credit Deposit Account No. **09-0457(IBM)** as described below.
- ☐ Charge the amount of _____
- ☒ Credit any overpayment.
- ☒ Charge any additional fee required.


Signature

Dated: **3/25/2004**

Jack P. Friedman
Reg. No. 44,688
Schmeiser, Olsen & Watts
3 Lear Jet Lane, Suite 201
Latham, NY 12110
(518) 220-1850

I certify that this document and fee is being deposited on **3/25/2004** with the U.S. Postal Service as first class mail under 37 C.F.R. 1.8 and is addressed to the Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.


Signature of Person Mailing Correspondence

Kim Dwileski

Typed or Printed Name of Person Mailing Correspondence

CC: